

# Deep Hierarchical Video Compression

Ming Lu<sup>1,2</sup>, Zhihao Duan<sup>3</sup>, Fengqing Zhu<sup>3</sup>, and Zhan Ma<sup>1\*</sup>

<sup>1</sup>School of Electronic Science and Engineering, Nanjing University

<sup>2</sup>Interdisciplinary Research Center for Future Intelligent Chips (Chip-X), Nanjing University

<sup>3</sup>Elmore Family School of Electrical and Computer Engineering, Purdue University  
minglu@nju.edu.cn, duan90@purdue.edu, zhu0@purdue.edu, mazhan@nju.edu.cn

## Abstract

Recently, probabilistic predictive coding that directly models the conditional distribution of latent features across successive frames for temporal redundancy removal has yielded promising results. Existing methods using a single-scale Variational AutoEncoder (VAE) must devise complex networks for conditional probability estimation in latent space, neglecting multiscale characteristics of video frames. Instead, this work proposes hierarchical probabilistic predictive coding, for which hierarchal VAEs are carefully designed to characterize multiscale latent features as a family of flexible priors and posteriors to predict the probabilities of future frames. Under such a hierarchical structure, lightweight networks are sufficient for prediction. The proposed method outperforms representative learned video compression models on common testing videos and demonstrates computational friendliness with much less memory footprint and faster encoding/decoding. Extensive experiments on adaptation to temporal patterns also indicate the better generalization of our hierarchical predictive mechanism. Furthermore, our solution is the first to enable progressive decoding that is favored in networked video applications with packet loss.

## Introduction

Deep learning breathes fresh life into the visual signal (e.g., images and videos) compression community that has been dominated by handcrafted codecs for decades (Wallace 1991; Marcellin et al. 2000; Wiegand et al. 2003; Sullivan et al. 2012; Bross et al. 2021). Instead of manually designing and optimizing individual modules such as transforms, mode selection, and quantization in traditional codecs, data-driven approaches adopt end-to-end learning of neural networks (Ballé, Laparra, and Simoncelli 2016; Theis et al. 2017). Despite the conceptual simplicity, learned image compression methods have achieved superior rate-distortion performance, surpassing the latest VVC (Versatial Video Coding (Bross et al. 2021)) intra codec (He et al. 2022; Lu et al. 2022).

For videos, however, learning-based methods are still not free from the shackles of the traditional hybrid framework. Most existing methods follow the two-stage pipeline shown

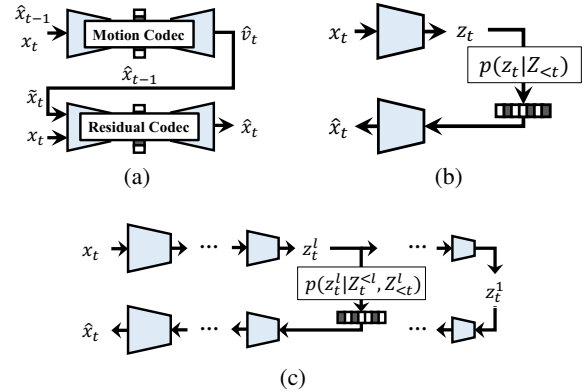


Figure 1: Interframe coding using (a) hybrid motion & residual coding, (b) single-scale probabilistic predictive coding, and (c) hierarchical probabilistic predictive coding (Ours).

in Fig. 1a: code motion flows first and then the residual between the current and motion-warped frame, either in an explicit (Lu et al. 2019) or conditional (Li, Li, and Lu 2021) manner. This framework is usually cumbersome in design (for example, separate models for intraframe coding, inter residual coding, motion coding, and motion estimation are required); thus, extensive hyperparameter tuning is necessary. Furthermore, inaccurate motion-induced warping error propagates inevitably across temporal frames, gradually degrading the quality of reconstructed frames over time.

As a promising solution to the problems mentioned earlier, (latent-space) probabilistic predictive coding attempts to reduce temporal redundancy by conditionally predicting future frames in a one-shot manner. Intuitively, if the current frame can be well predicted through the past frames, motion (e.g., flow) estimation and compensation can be completely exempted, and the aforementioned error propagation can also be eliminated accordingly. Recently, Mentzer *et al.* (Mentzer et al. 2022) proposed a probabilistic predictive video coding framework named *Video Compression Transformer (VCT)*. Under the VAE-based image compression framework, VCT models the latent features of the current frame conditioned on the previous-frame latent features using a transformer-based temporal entropy model. Though VCT outperforms many previous video coding methods, its

\*Corresponding Author

conditional prediction of single-scale latent features at 1/16 resolution of the original frame in Fig. 1b fundamentally constrains its characterization capacity, which ignores multiscale characteristics of video frames.

This paper proposes a hierarchical probabilistic predictive coding, termed DHVC, in which conditional probabilities of multiscale latent features of future frames are effectively modeled using deliberately-designed, powerful hierarchical VAEs. The latent distribution at a certain scale in the current frame is predicted by the prior features from previous scales in the same frame and the corresponding scale of the previous frames. Doing so gives us a powerful and efficient modeling ability to characterize arbitrary feature distributions. For instance, Mentzer *et al.* (Mentzer et al. 2022) relied on a complicated prediction in a block-level autoregressive manner, which is inefficient. Instead, we perform a multi-stage conditional probability prediction, showing better performance and desiring less complexity.

Upon extensive evaluations using commonly used video sequences, our method outperforms well-known learned models using hybrid motion and residual coding and previous state-of-the-art method using latent probabilistic predictive coding. Extensive studies on the adaptation to various temporal patterns also reveal the generalization of our hierarchical predictive mechanism. In addition, our method also supports temporal progressive decoding, being the first learned progressive video coding method to our best knowledge. Therefore, it can handle packet losses induced by poor network connections to some extent.

Our contributions can be summarized as follows:

- We propose a hierarchical probabilistic prediction model for video coding. Our model employs a collection of multiscale latent variables to represent the coarse-to-fine nature of video frames scale-wisely.
- We propose the spatial-temporal prediction and in-loop decoding fusion modules, which enable better performance, lower memory consumption, and faster encoding/decoding than the previous best probabilistic predictive coding-based method (Mentzer et al. 2022).
- Experiments demonstrate that our method is better generalized to various temporal patterns. Our model is also the first to support the functionality of progressive decoding.

## Related Work

We briefly review end-to-end learned video coding methods, including classical hybrid motion and residual coding and recently-emerged probabilistic predictive coding. We also theoretically explain the hierarchical VAE formalism as it provides the basis for our method.

### Learned Video Coding

**Data compression and variational autoencoders (VAEs):** Let  $x$  denote data (e.g., image or video) with an unknown distribution. Traditional image/video coding belongs to the *transform coding*, where one wants to find an encoder  $f_e$ , a decoder  $f_d$ , and an entropy model for the transform coefficients such that the rate-distortion cost is minimized:

$$\min H(f_e(x)) + \lambda \cdot d(x, f_d(f_e(x))). \quad (1)$$

Here, the first term is the (cross-) entropy of the compressed coefficients approximating the rate,  $d$  is a distortion function, and  $\lambda$  is the Lagrange multiplier that balances the rate and distortion tradeoff. As studied in (Ballé et al. 2018; Duan et al. 2023), transform coding can be equivalently considered as data distribution modeling using variational autoencoders (VAEs). Specifically, VAEs assume a model of data:

$$p(x, z) = p(x | z) \cdot p(z), \quad (2)$$

where  $z$  is latent variables like transform coefficients. In VAE, a *prior*  $p(z)$  describes the distribution of the latent variable, a decoder  $p(x | z)$  maps latent-space elements to original data-space signal, and an *approximate posterior*  $q(z | x)$  (i.e., the encoder) encodes data into the latent space. Letting  $\hat{x} \sim p(x | z)$  denote the reconstruction, the objective can be written as (Yang and Mandt 2022; Duan et al. 2023)

$$\min D_{\text{KL}}(q(z | x) \| p(z)) + \lambda \cdot d(x, \hat{x}), \quad (3)$$

and if the posterior  $q(z | x)$  is deterministic and discrete (e.g., when quantization is applied to  $z$ ), this VAE objective equals the rate-distortion optimization in Eq. (1). Such a connection has inspired many subsequent works to apply VAE-based probabilistic methods to compression task, such as (Yang, Bamler, and Mandt 2020b; Agustsson and Theis 2020; Yang, Bamler, and Mandt 2020a; Theis and Ahmed 2022; Ryder et al. 2022; Chen et al. 2022).

Learned video coding methods can be generally categorized into two groups: hybrid motion & residual coding and probabilistic predictive coding.

**Hybrid Motion & Residual Coding** refers to the classical coding framework with motion and residual processing. Lu et al. (Lu et al. 2019) first proposed to use two similar VAE-based networks to code the motion and residuals, respectively, which was then enhanced with better motion alignment in (Lu et al. 2020; Liu et al. 2020a). Then, Hu et al. (Hu, Lu, and Xu 2021) migrated the motion alignment to the feature domain and achieved better compression performance. Recently, by converting residual coding to conditional coding of aligned features, Li et al. (Li, Li, and Lu 2021) took the learned video coding to a new level of performance. Subsequently, by further integrating multi-scale aligned feature fusion, post-processing, and bitrate allocation, learned video coding algorithms achieved unprecedented compression efficiency, surpassing the latest VVC (Li, Li, and Lu 2022).

**Probabilistic Predictive Coding** is an emerging video coding method. Liu et al. (Liu et al. 2020b) relied on stacked convolutions for latent distribution prediction, while VCT (Mentzer et al. 2022) adopted Transformer for the same purpose. Both works perform temporally-conditional distribution prediction only using single-scale latent variables (i.e., 1/16 of the original resolution), which greatly constrains the accuracy of probability estimation and leads to sub-optimal predictive performance. Therefore, in this paper, we propose a hierarchical probabilistic predictive coding method, which substantially improves the accuracy and efficiency of temporal prediction by characterizing multiscale latent features for conditional probability estimation in a coarse-to-fine approach.

## Hierarchical VAEs

To improve the flexibility and expressiveness of single-scale VAE, hierarchical VAEs (Kingma et al. 2016; Child 2020; Vahdat and Kautz 2020) employ multiscale latent variables, denoted by  $Z = \{z^1, \dots, z^L\}$ . Accordingly, the latent priors can be factorized as:

$$p(Z) = \prod_{l=1}^L p(z^l | Z^{<l}), \quad (4)$$

where  $L$  is the total number of hierarchical scales, and  $Z^{<l}$  denotes  $\{z^1, \dots, z^{l-1}\}$ . Typically,  $z^1$  has the smallest dimension, while  $z^L$  is with the largest dimension. Such a dimensional refinement from a lower scale to a higher one improves the flexibility of VAEs and effectively captures the coarse-to-fine characteristics of images.

Among popular hierarchical VAE architectures, ResNet VAE (Kingma et al. 2016) provides the most promising performance in terms of image modeling. Different from the Hyperprior VAE used in (Ballé et al. 2018), each latent variable  $z^l$  of ResNet VAE is conditioned on all  $Z^{<l}$  and its encoding is bi-directional with dependent on both  $x$  and  $Z^{<l}$ . This might explain the fact that ResNet VAE can be scaled up to more than 70 layers (Child 2020). The loss function for training ResNet VAEs can be extended from Eq. (3) for supervising multiscale latents:

$$\min \sum_{i=1}^L D_{\text{KL}}(q^i \parallel p^i) + \lambda \cdot d(x, \hat{x}), \quad (5)$$

where  $q^l$  and  $p^l$  are shorthand notations for the posterior and prior for the  $l$ -th scale latent variable, i.e.,

$$q^l = q(z^l | x, Z^{<l}), \quad \text{and} \quad p^l = p(z^l | Z^{<l}). \quad (6)$$

Our proposed method is developed based on the ResNet VAE structure by further introducing the temporal priors in addition to the deliberate design for practical compression.

## Preliminary: Predictive Video Coding

Suppose a video sequence  $X = \{x_1, \dots, x_T\}$  that contains  $T$  frames for encoding. As a convention for predictive coding methodologies (Liu et al. 2020b; Mentzer et al. 2022), an analysis encoder followed by quantization is first applied to transform each input frame  $x_t$  into the discrete latent representation  $z_t$  with reduced resolution. A symmetrical decoder is then used to recover the reconstruction  $\hat{x}_t$  from  $z_t$ . Having the probability mass function (PMF)  $p(z_t)$  to estimate the true distribution of symbols in  $z_t$ , we can get the bits desired for transmission by entropy coding. The main idea of the probabilistic predictive coding is to parameterize the  $p(z_t)$  as a conditional distribution

$$p(z_t) = p(z_t | Z_{<t}), \quad (7)$$

where  $Z_{<t}$  is a set of latent features preceding time  $t$ . By exploiting the temporal redundancy across frames using an efficient prediction network, one can obtain more accurate probability estimation for the current frame to reduce cross-entropy and thus maximize coding efficiency. To this end,

VCT (Mentzer et al. 2022) introduces a Transformer model in a block-level autoregressive manner to model  $p(z_t^i | z_t^{<i}, Z_{<t})$  for jointly appreciating spatial and temporal correlation. Here,  $z_t^i$  corresponds to a pixel at position  $i$  within a predefined block in the latent space of the current frame, and  $z_t^{<i}$ s are autoregressive neighbors previously processed.

Although decent compression performance is achieved, VCT only performs conditional probability prediction using single-scale latent variables, ignoring the multiscale characteristics in both spatial and temporal domains. Therefore, the prediction is sub-optimal, and the complexity of the prediction network is usually unaffordable.

## Proposed Method

### Network Architecture Overview

Figure 2a depicts the overall framework of our method, which consists of a *bottom-up* path and a *top-down* path. Given an input frame  $x_t$ , the bottom-up path produces a set of features  $R_t = \{r_t^1, \dots, r_t^L\}$  at respective 1/64, 1/32, 1/16, and 1/8 resolutions of the original input through downsampling and feature aggregation/embedding using residual blocks (ResBlocks). Different from (Chen et al. 2022), the  $r_t^l$  at each scale is dependent on both feature extracted from  $x_t$  and decoded feature from lower scale, which turns our method into conditional residual coding to minimize the bitrate consumption. These residual features  $R_t$  are subsequently sent to the top-down path for hierarchical probabilistic modeling.

The top-down path starts from two learnable constant biases and then encodes a sequence of latent variables  $Z_t = \{z_t^1, \dots, z_t^L\}$  (in the Latent Blocks) to produce respective prior feature  $f_t^l$  and reconstructive feature  $d_t^l$  scale-by-scale. In the end, the  $\hat{x}_t$  is reconstructed by passing the last reconstructive representation  $d_t^L$  into multiple upsampling and ResBlock layers. The down-sampling operations  $\downarrow$  are implemented by strided convolution, and the up-sampling operations  $\uparrow$  are implemented by  $1 \times 1$  convolution followed by pixel shuffle. The ConvNeXt (Liu et al. 2022) units are adopted in ResBlocks. More details can be found in supplementary materials.

### Predictive Coding Modules

We now detail the architecture of the Latent Block (Fig. 2b), which is critical to the effectiveness of our approach. Like in ResNet VAEs, each Latent Block adds “information”, carried by the latent variable  $z_t^l$ , into the top-down path features. We substantially extend it by introducing (1) a Spatial-Temporal Prediction module for predictive coding and (2) an In-loop Decoding Fusion module to improve coding performance, which is described below one-by-one.

**Spatial-Temporal Prediction** (Fig. 2c): To predict  $z_t^l$  at  $l$ -th scale, we combine the same-scale temporal priors  $Z_{<t}^l$  with spatial prior  $f_t^{l-1}$  from previous scales to produce the prior distribution parameters. We begin with the Temporal Fusion by passing the temporal priors  $Z_{<t}^l$  into stacked ResBlocks, with skip connections at each level. Then, the spatial prior feature  $f_t^{l-1}$  is concatenated with the fused temporal information for subsequent Conditional Generation to

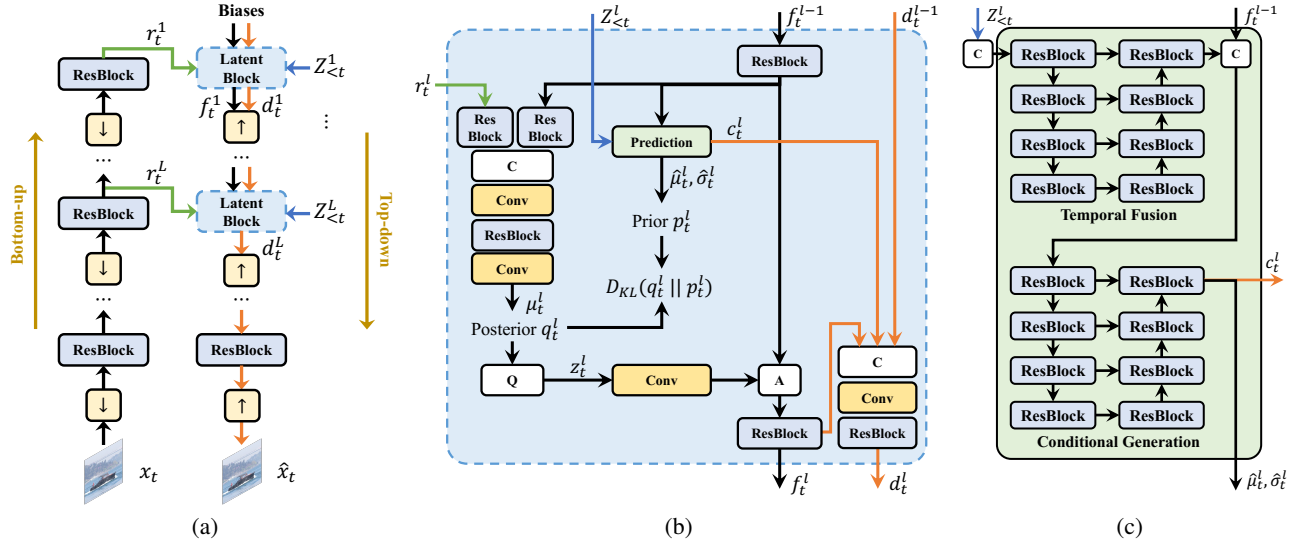


Figure 2: (a) Overall Architecture, (b) Latent Block, and (c) Spatial-Temporal Prediction Module of our proposed DHVC.  $\downarrow$  and  $\uparrow$  are respective downscaling and upscaling operations. “C” represents concatenation, “A” represents addition, and “Q” represents quantization. The convolutional layer “Conv” is used for feature re-dimension.

get the contextual feature  $c_t^l$  and the prior distribution parameters, i.e., the mean  $\hat{\mu}_t^l$  and scale  $\hat{\sigma}_t^l$ .

**In-loop Decoding Fusion** (on the right of Fig. 2b): Two distinct features are generated during the decoding process: the prior feature  $f_t^l$  utilized as the spatial prior for subsequent scale, and the reconstructive feature  $d_t^l$  for the eventual reconstruction. In our specific implementation, we concatenate the previously decoded feature  $d_t^{l-1}$  and the contextual feature  $c_t^l$ , along with the  $f_t^l$  to generate the fused results  $d_t^l$ . This design represents a notable departure from the original ResNet VAE framework, which employs a single top-down path feature for both the prior and reconstruction purposes. Through the method in our study, the  $f_t^l$  solely handles conditional distribution modeling, whereas the  $d_t^l$  is responsible for the reconstruction aspect. By leveraging the dependable contextual feature  $c_t^l$ , we achieve a desirable decoded  $d_t^l$  while conserving bitrate consumption effectively.

Note that our framework requires only a single model for intra and inter frame coding. For intra-coding, the temporal priors  $Z_{<t}^l$  each scale are set using learnable constant biases, while for inter coding, we use the two previous latents  $z_{t-1}^l$  and  $z_{t-2}^l$  to make up of  $Z_{<t}^l$ .

### Probabilistic Model and Loss Function

With the specific neural network modules, our framework effectively extends hierarchical VAEs to predictive video coding. To support practical lossy compression using feasible entropy coding algorithms, we follow previous works (Ballé et al. 2018; Duan et al. 2023) and apply quantization-aware training using uniform posteriors. Specifically, we adopt a hybrid quantization strategy at training time to simulate the quantization error. The additive uniform noise is applied in terms of rate estimation, while the

straight-through rounding operation is used for reconstruction. We use uniform quantization at test time. For the prior, we use the Gaussian distribution convolved with uniform distribution, which is flexible to match the posterior.

**Posteriors:** The (approximate) posterior for the  $l$ -th latent variable,  $z_t^l$ , is defined to be an uniform distribution:

$$p(z_t^l | x_t, Z_{<t}^{<l}) = U(\mu_t^l - \frac{1}{2}, \mu_t^l + \frac{1}{2}), \quad (8)$$

where  $\mu_t^l$  is the output of the posterior branch in the latent block (see Fig. 2b) by merging the embedded feature  $r_t^l$  and prior feature  $f_t^{l-1}$  from the previous level. The discrete  $z_t^l$  depends on the frame  $x_t$  as well as previous level latent variables  $Z_{<t}^{<l}$ .

**Priors:** Our framework extends ResNet VAE to predictive video coding by conditioning the prior distributions for  $z_t^l$  on temporal latent variables  $Z_{<t}^l$ . At each timestep, considering  $L$  levels of latent variables  $Z_t = \{z_t^1, \dots, z_t^L\}$ , the latent conditional distribution can be factorized as

$$p(Z_t | Z_{<t}) = \prod_{l=1}^L p(z_t^l | Z_{<t}^{<l}, Z_{<t}^l), \quad (9)$$

Then, the prior distribution for each  $z_t^l$  is defined as a Gaussian convolved with a uniform distribution:

$$p(z_t^l | Z_{<t}^{<l}, Z_{<t}^l) = \mathcal{N}(\hat{\mu}_t^l, \hat{\sigma}_t^{l^2}) * U(-\frac{1}{2}, \frac{1}{2}) \quad (10)$$

where  $\mathcal{N}(\hat{\mu}_t^l, \hat{\sigma}_t^{l^2})$  denotes the Gaussian probability density function. The mean  $\hat{\mu}_t^l$  and scale  $\hat{\sigma}_t^l$  are predicted by the prior branch in the latent block. Note that the prior mean  $\hat{\mu}_t^l$  and scale  $\hat{\sigma}_t^l$  are dependent on both the latent variables from previous time steps  $Z_{<t}^l$  for that level and on the latent variables of the previous levels at the current timestep  $Z_{<t}^{<l}$ .

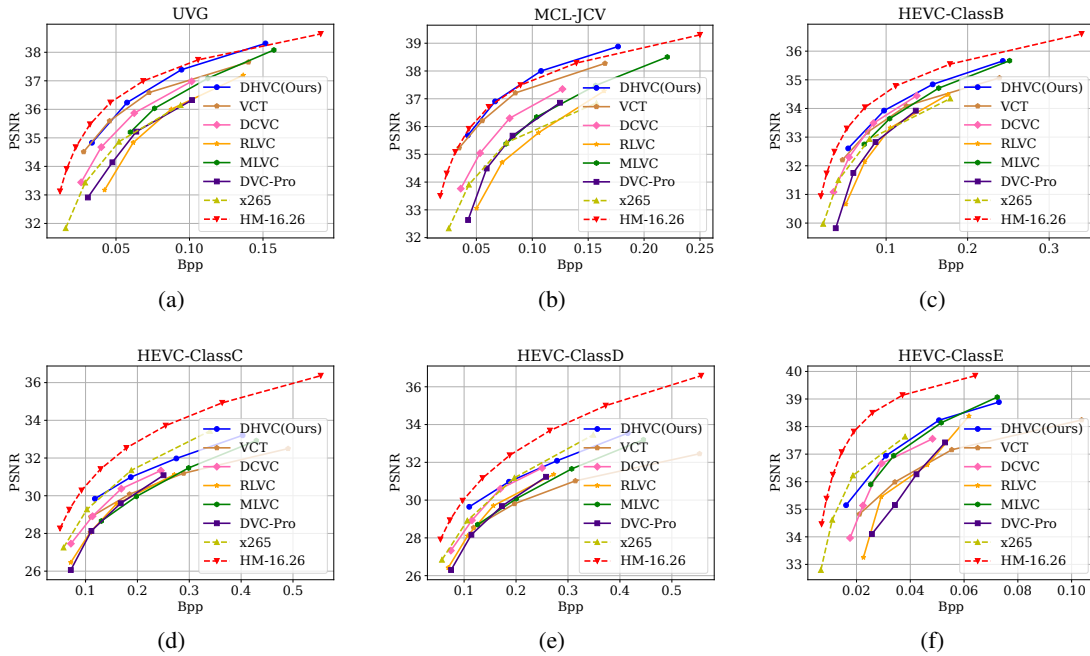


Figure 3: Compression efficiency comparison using rate-distortion (R-D) curves.

**Training Objective:** Typically, the hybrid motion and residual coding methods require multi-stage or simultaneous optimization of the optical flow, motion coding, and residual coding networks during the training phase. Differently, the training process of our model is as easy as optimizing a lossy image coding. The loss function  $\mathcal{L}$  is extended from Eq. (5) with the inclusion of temporal dependency:

$$\mathcal{L} = \min \sum_{l=1}^L -\log_2 p(z_t^l | Z_t^{<l}, Z_{<t}^l) + \lambda \cdot d(x_t, \hat{x}_t), \quad (11)$$

The first term consists of the rate for all latent variables. The second term corresponds to the reconstruction distortion, which is commonly chosen to be the Mean Squared Error (MSE) or MS-SSIM (Wang, Simoncelli, and Bovik 2003) loss for videos. The multiplier  $\lambda$ , which trades off rate and distortion, is pre-determined and fixed throughout training. At test time, the  $\lambda$  is the same for both intra and inter frame, and the actual bitrates are determined by the estimation accuracy of conditional probabilistic modeling.

## Experimental Results

### Implementation Settings

**Datasets:** We use the popular Vimeo-90K (Xue et al. 2019) dataset to train our model, which consists of 64,612 video samples. Training batches comprise sequential frames that are randomly cropped to the size of  $256 \times 256$ . Commonly used test datasets, i.e., the UVG (Mercat, Viitanen, and Vanne 2020), MCL-JCV (Wang et al. 2016), and HEVC Class B, C, D, and E (Bossen et al. 2013), are used for evaluation. They cover various scene variations, and resolutions are available from  $416 \times 240$  to  $1920 \times 1080$ . Test sequences

in YUV420 format are pre-processed following the suggestions in (Sheng et al. 2022) to generate RGB frames as the input of learned models. The first 96 frames of each video are used for evaluation, and the group of pictures (GOP) is set at 32. These settings are the same as in other learned video coding methods.

**Training details:** We progressively train our model for fast convergence. First, the model is trained to encode a single frame independently for 2M iterations by setting the temporal prior at each scale level as a learnable bias. Then, we train the aforementioned model for 500K steps using three successive frames, with temporal priors hierarchically generated from previously-decoded frames. In the end, another 100K steps are applied to fine-tune the model using five successive frames, for which it better captures long-term temporal dependence (Liu et al. 2020a). We set  $\lambda$  from  $\{256, 512, 1024, 2048\}$  and  $\{4, 8, 16, 32\}$  for respective MSE and MS-SSIM optimized models to cover wide rate ranges. Adam (Kingma and Ba 2014) is the optimizer with the learning rate at  $10^{-4}$ . Our model is trained using two Nvidia RTX 3090, and the batch size is fixed at eight.

### Evaluation

All the evaluation experiments are performed under the low-delay configuration. We choose x265<sup>1</sup> and HM-16.26<sup>2</sup> as the benchmarks of traditional video codecs. Both of them use the default configuration. The detailed codec settings can be found in supplementary materials. For learned video coding methods, we compare with the representative algorithms using hybrid motion & residual coding

<sup>1</sup><https://www.videolan.org/developers/x265.html>

<sup>2</sup><https://hevc.hhi.fraunhofer.de>

Method	Parm (M)	Ops	Mem (G)	ET/DT (s)
DCVC	<b>7.58</b>	1126.40	11.71	13.78/46.48
VCT	187.82	3042.20	10.51	1.64/1.58
DHVC	112.46	<b>433.81</b>	<b>4.27</b>	<b>0.25/0.21</b>

Table 1: Complexity comparison among DCVC, VCT, and DHVC (Ours). “Parm” represents the model parameters. “Ops” denotes the average multiply-add operations per pixel. “Mem” is the peak memory consumed in the inference. “ET” and “DT” are the encoding and decoding times. We perform evaluations on a single RTX 3090-24G GPU.

method, i.e., DVC-Pro (Lu et al. 2020), MLVC (Lin et al. 2020), RLVC (Yang et al. 2020), and DCVC (Li, Li, and Lu 2021), and the best-performing probabilistic predictive coding model VCT (Mentzer et al. 2022). Recently, a set of codec optimization tools, like feature fusion, post-processing, bitrate allocation, etc., are integrated into the representative frameworks to further improve the compression performance (Li, Li, and Lu 2022). In this work, we currently focus on comparisons regarding frameworks themselves, and choose the best-in-class methods without further augmentation. Integrating these optimization mechanisms into our method is an interesting topic for future study.

**Performance:** Figure 3 depicts the R-D (rate-distortion) curves for testing various methods across various datasets. Our model, DHVC, leads all the learned methods regardless of testing videos, revealing the generalization of our method. It also performs much better than x265 on popular datasets with 1080p resolution (e.g., UVG and MCL-JCV) and even better than HM-16.26 on the MCL-JCV, suggesting the encouraging potential of such hierarchical predictive coding.

Though our method still performs the best, we must admit that it yet remains a noticeable performance gap between the learned codecs and traditional ones on low-resolution videos in HEVC Class C, D, and E. We believe this mainly owes to the conditional probability estimation or motion/residual coding upon latent variable with much lower resolution downsampled from the original input. For a low-resolution input, its local block contains more complicated texture patterns than that in a same-size block of a high-resolution video. Thus, the downsampling-induced information loss is more critical, potentially leading to inaccurate correlation characterization and subsequent compression. Such a hypothesis is well justified when comparing the performance between the VCT and our method. VCT shows a great performance drop for those low-resolution datasets, even worse than the earliest learned method DVC-Pro in middle and high bitrate ranges. This is because VCT conducts the conditional probability estimation upon single-scale latent features at 1/16 resolution of the original input. Instead, our method performs conditional probability estimation using multi-scale latent variables at respective 1/64, 1/32, 1/16, and 1/8 resolutions. Such a hierarchical mechanism greatly improves performance by thoroughly exploiting the coarse-to-fine natural characteristics of video frames.

Due to the space limitation, results for learned models




Figure 4: Performance contribution of modular components.

trained using MS-SSIM loss can be found in the supplemental materials, which show a clear advantage of our method over both the traditional and learned video codecs.

**Complexity:** Evaluation results with 1080p videos are listed in Table 1. Except for the model size, our method shows clear advantages for other metrics, reporting the least requirements of respective kMACs per pixel, peak memory consumption, encoding, and decoding time. This also suggests that the model size is not closely related to the computational complexity of running codecs in practice. The sizeable parameters used in our method are mainly attributed to using basic ConvNeXt units to form the ResBlocks and Latent Blocks (see Fig. 2a).

Our DHVC shows a clear reduction in kMACs per pixel and peak memory occupation, owing to the use of simple probabilistic prediction modules instead of complicated Transformer-based prediction network in VCT or complicated motion and residual coding modules in DCVC.

For encoding and decoding time, as the DCVC applies the pixel-wise spatial autoregressive model for entropy coding, it takes about 17.86 and 40.64 seconds, which is unacceptable for practical codecs. VCT, instead, uses a simplified  $4 \times 4$  block-level spatial autoregressive model, offering faster encoding (decoding) to DCVC. Our method completely removes the use of a spatial autoregressive model through the proposed hierarchical processing pipeline, which further reduces the encoding and decoding time to respective 0.25 and 0.21 seconds, i.e.,  $55 \times / 221 \times (6 \times / 7 \times)$  faster encoding/decoding than the DCVC and VCT respectively.

## Deep Dive

We perform ablation studies to understand the capacity of our proposed DHVC better.

**Modular Contribution:** We further examine the contribution of each module in our proposed DHVC in Fig. 4. “Baseline” denotes the model disabling both the temporal prediction and in-loop decoding fusion in latent blocks, with only the spatial prior from previous scales for probabilistic modeling. “Baseline + TP” indicates the temporal probabilistic prediction is integrated to reduce the temporal redundancy. Apparently, the performance with the support of temporal information improves significantly upon the base model. Furthermore, with the help of in-loop decoding fusion module, dubbed by “Baseline + TP + DF” in the figure, an averaged 1 dB PSNR increase is obtained. It is sufficient to justify the advantage of compensating for fusion based on

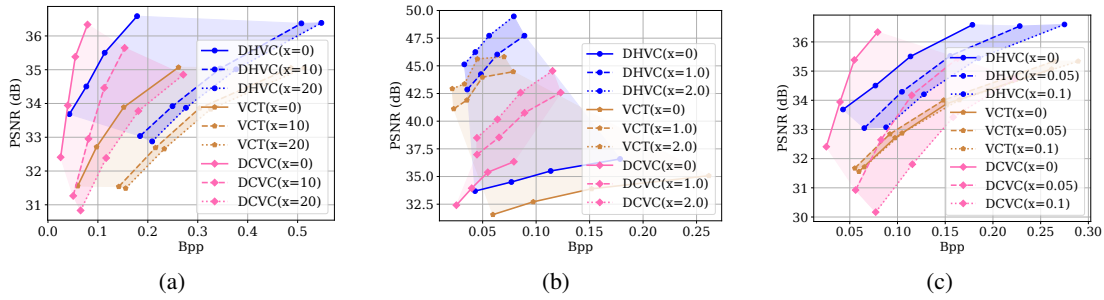


Figure 5: Impact of temporal pattern on compression using synthetic data: (a) pixel shifting with values  $x = 0, 10, 20$ , (b) Gaussian blurring with sigma  $x \cdot t$  at frame order  $t$ , and (c) fading by linear transition between two unrelated scenes using alpha blending. All evaluations start from  $x = 0$  (i.e., videos consist of still images) as denoted by solid lines in the figures.




Figure 6: Progressive decoding by visualizing five scale levels. Blue bars and red lines represent bit per pixel (bpp) and PSNR values for each frame respectively. Please zoom in for more details.

prediction results on the decoding side. In addition, the long-term finetuning with five frames, represented by “Baseline + TP +DF + LT”, brings another R-D improvement to constitute the complete performance of our work. This suggests that the rate-distortion relationship between frames can be effectively balanced by joint training with multiple frames.

**Adaptation Capacity to Temporal Patterns** is critical for the model’s generalization when encoding different video contents. We exactly follow (Mentzer et al. 2022) to generate videos using three different temporal patterns, including the pixel shifting, blurring, and fading effects. The R-D curves are plotted in Fig 5. Our method outperforms the VCT with regard to all synthetic datasets, which demonstrates the powerful modeling capability of the hierarchical probabilistic predictive mechanism. No matter which temporal pattern or how fast the scene change is, our method is consistently applicable. However, both the VCT and our DHVC behave worse than the DCVC when it comes to the videos with pixel-shifted. This is mainly due to the fact that the DCVC utilizes a motion alignment module by encoding the motion data. For such regular object displacement, motion estimation can achieve high prediction accuracy, which has obvious advantages over our latent space probabilistic prediction. Considering how to add hierarchical motion alignment to our approach is a topic worth future exploring.

**Progressive Decoding Capability** is enabled in the proposed DHVC, which was seldom supported in existing methods. Specifically, once obtain the lowest-scale features (lowest resolution) of the current frame, we have the coars-

est frame reconstruction after decoding as in the upper left subplot of Fig. 6 (e.g., Level 1). As additional compressed latent features are transmitted to the decoder side, we can clearly observe the improvement of reconstruction results (see visualized subplots with more scale levels and PSNR increases in bottom subplots accordingly). When we receive partial scales, we notice the PSNR degradation (red curves) in a GOP. This is due to the error propagation since the temporal references can only provide partial priors for decoding. Instead, once we receive all-scale latent features, the PSNR metric is stable across frames and GOPs (see “Level 1 - 5”). At the same time, gradual PSNR degradation in a GOP still prevails in existing video codecs using hybrid motion & residual coding.

Progressive decoding quickly provides relatively-coarse reconstructions by encoding and transmitting partial features. In this exemplified Fig. 6, for such a 1080p video, having two scales of compressed latent features can already present a clear preview of the content, by which our model provides a fast and less bitrate-consuming preview in video streaming applications. This also gives us a broader understanding: we can still decode the content in networked applications with packet loss when we have partial packets. In congested connections, we can proactively drop latent packets corresponding to higher scales.

## Conclusion

This paper proposes a novel hierarchical probabilistic predictive coding framework for learning-based video compression, termed DHVC. The DHVC provides superior compression efficiency to popular and representative learned video codecs across a great variety of video samples. More importantly, DHVC offers the fastest encoding and decoding with the least running memory, which not only reveals the best balance between coding performance and complexity efficacy but also offers encouraging potential for application of learned video codecs. Our future work will focus on exploring efficient prior representations or optimization mechanisms to further improve the compression efficiency.

## Acknowledgements

This work is partially supported by the National Key Research and Development Project of China (No.2022YFF0902402) and the Natural Science Foundation of China (No.U20A20184).

## References

- Agustsson, E.; and Theis, L. 2020. Universally Quantized Neural Compression. *Advances in Neural Information Processing Systems*, 33: 12367–12376.
- Ballé, J.; Laparra, V.; and Simoncelli, E. P. 2016. End-to-end optimized image compression. *arXiv preprint arXiv:1611.01704*.
- Ballé, J.; Minnen, D.; Singh, S.; Hwang, S. J.; and Johnston, N. 2018. Variational image compression with a scale hyperprior. *arXiv preprint arXiv:1802.01436*.
- Bossen, F.; et al. 2013. Common test conditions and software reference configurations. *JCTVC-L1100*, 12(7): 1.
- Bross, B.; Wang, Y.-K.; Ye, Y.; Liu, S.; Chen, J.; Sullivan, G. J.; and Ohm, J.-R. 2021. Overview of the versatile video coding (VVC) standard and its applications. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(10): 3736–3764.
- Chen, Z.; Gu, S.; Lu, G.; and Xu, D. 2022. Exploiting intra-slice and inter-slice redundancy for learning-based lossless volumetric image compression. *IEEE Transactions on Image Processing*, 31: 1697–1707.
- Child, R. 2020. Very deep vaes generalize autoregressive models and can outperform them on images. *arXiv preprint arXiv:2011.10650*.
- Duan, Z.; Lu, M.; Ma, Z.; and Zhu, F. 2023. Lossy Image Compression with Quantized Hierarchical VAEs. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 198–207.
- He, D.; Yang, Z.; Peng, W.; Ma, R.; Qin, H.; and Wang, Y. 2022. Elic: Efficient learned image compression with unevenly grouped space-channel contextual adaptive coding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5718–5727.
- Hu, Z.; Lu, G.; and Xu, D. 2021. FVC: A new framework towards deep video compression in feature space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1502–1511.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kingma, D. P.; Salimans, T.; Jozefowicz, R.; Chen, X.; Sutskever, I.; and Welling, M. 2016. Improved variational inference with inverse autoregressive flow. *Advances in neural information processing systems*, 29.
- Li, J.; Li, B.; and Lu, Y. 2021. Deep contextual video compression. *Advances in Neural Information Processing Systems*, 34: 18114–18125.
- Li, J.; Li, B.; and Lu, Y. 2022. Hybrid spatial-temporal entropy modelling for neural video compression. In *Proceedings of the 30th ACM International Conference on Multimedia*, 1503–1511.
- Lin, J.; Liu, D.; Li, H.; and Wu, F. 2020. M-LVC: Multiple frames prediction for learned video compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3546–3554.
- Liu, H.; Lu, M.; Ma, Z.; Wang, F.; Xie, Z.; Cao, X.; and Wang, Y. 2020a. Neural video coding using multiscale motion compensation and spatiotemporal context model. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(8): 3182–3196.
- Liu, J.; Wang, S.; Ma, W.-C.; Shah, M.; Hu, R.; Dhawan, P.; and Urtasun, R. 2020b. Conditional entropy coding for efficient video compression. In *European Conference on Computer Vision*, 453–468. Springer.
- Liu, Z.; Mao, H.; Wu, C.-Y.; Feichtenhofer, C.; Darrell, T.; and Xie, S. 2022. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11976–11986.
- Lu, G.; Ouyang, W.; Xu, D.; Zhang, X.; Cai, C.; and Gao, Z. 2019. Dvc: An end-to-end deep video compression framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11006–11015.
- Lu, G.; Zhang, X.; Ouyang, W.; Chen, L.; Gao, Z.; and Xu, D. 2020. An end-to-end learning framework for video compression. *IEEE transactions on pattern analysis and machine intelligence*, 43(10): 3292–3308.
- Lu, M.; Guo, P.; Shi, H.; Cao, C.; and Ma, Z. 2022. Transformer-based Image Compression. In *2022 Data Compression Conference (DCC)*, 469–469.
- Marcellin, M. W.; Gormish, M. J.; Bilgin, A.; and Boliek, M. P. 2000. An overview of JPEG-2000. In *Proceedings DCC 2000. Data Compression Conference*, 523–541. IEEE.
- Mentzer, F.; Toderici, G.; Minnen, D.; Hwang, S.-J.; Caelles, S.; Lucic, M.; and Agustsson, E. 2022. Vct: A video compression transformer. *arXiv preprint arXiv:2206.07307*.
- Mercat, A.; Viitanen, M.; and Vanne, J. 2020. UVG dataset: 50/120fps 4K sequences for video codec analysis and development. In *Proceedings of the 11th ACM Multimedia Systems Conference*, 297–302.
- Ryder, T.; Zhang, C.; Kang, N.; and Zhang, S. 2022. Split Hierarchical Variational Compression. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 386–395.
- Sheng, X.; Li, J.; Li, B.; Li, L.; Liu, D.; and Lu, Y. 2022. Temporal Context Mining for Learned Video Compression. *IEEE Transactions on Multimedia*, 1–12.
- Sullivan, G. J.; Ohm, J.-R.; Han, W.-J.; and Wiegand, T. 2012. Overview of the high efficiency video coding (HEVC) standard. *IEEE Transactions on circuits and systems for video technology*, 22(12): 1649–1668.
- Theis, L.; and Ahmed, N. Y. 2022. Algorithms for the Communication of Samples. *Proceedings of the International Conference on Machine Learning*, 162: 21308–21328.
- Theis, L.; Shi, W.; Cunningham, A.; and Huszár, F. 2017. Lossy Image Compression with Compressive Autoencoders. *International Conference on Learning Representations*.



- Vahdat, A.; and Kautz, J. 2020. NVAE: A deep hierarchical variational autoencoder. *Advances in neural information processing systems*, 33: 19667–19679.
- Wallace, G. K. 1991. The JPEG still picture compression standard. *Communications of the ACM*, 34(4): 30–44.
- Wang, H.; Gan, W.; Hu, S.; Lin, J. Y.; Jin, L.; Song, L.; Wang, P.; Katsavounidis, I.; Aaron, A.; and Kuo, C.-C. J. 2016. MCL-JCV: a JND-based H. 264/AVC video quality assessment dataset. In *2016 IEEE international conference on image processing (ICIP)*, 1509–1513. IEEE.
- Wang, Z.; Simoncelli, E. P.; and Bovik, A. C. 2003. Multi-scale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, 1398–1402. Ieee.
- Wiegand, T.; Sullivan, G. J.; Bjontegaard, G.; and Luthra, A. 2003. Overview of the H. 264/AVC video coding standard. *IEEE Transactions on circuits and systems for video technology*, 13(7): 560–576.
- Xue, T.; Chen, B.; Wu, J.; Wei, D.; and Freeman, W. T. 2019. Video enhancement with task-oriented flow. *International Journal of Computer Vision*, 127: 1106–1125.
- Yang, R.; Mentzer, F.; Van Gool, L.; and Timofte, R. 2020. Learning for video compression with recurrent auto-encoder and recurrent probability model. *IEEE Journal of Selected Topics in Signal Processing*, 15(2): 388–401.
- Yang, Y.; Bamler, R.; and Mandt, S. 2020a. Improving inference for neural image compression. *Advances in Neural Information Processing Systems*, 33: 573–584.
- Yang, Y.; Bamler, R.; and Mandt, S. 2020b. Variational Bayesian Quantization. *Proceedings of the International Conference on Machine Learning*, 119: 10670–10680.
- Yang, Y.; and Mandt, S. 2022. Towards Empirical Sandwich Bounds on the Rate-Distortion Function. *International Conference on Learning Representations*.